2017 accounting issues: CMS

J. Flix, J. Letts

25/Jan./2018

What?

- We've found inconsistencies for some Tier-2 sites when producing 2017 resource utilization plots:
  - Data gathered from: https://accounting.egi.eu/
  - CPU utilization (walltime HS06·days) apparently fine… But low CPU eff by the end of 2017
  - We noticed some big Tier-2 sites were having accounting problems, since months

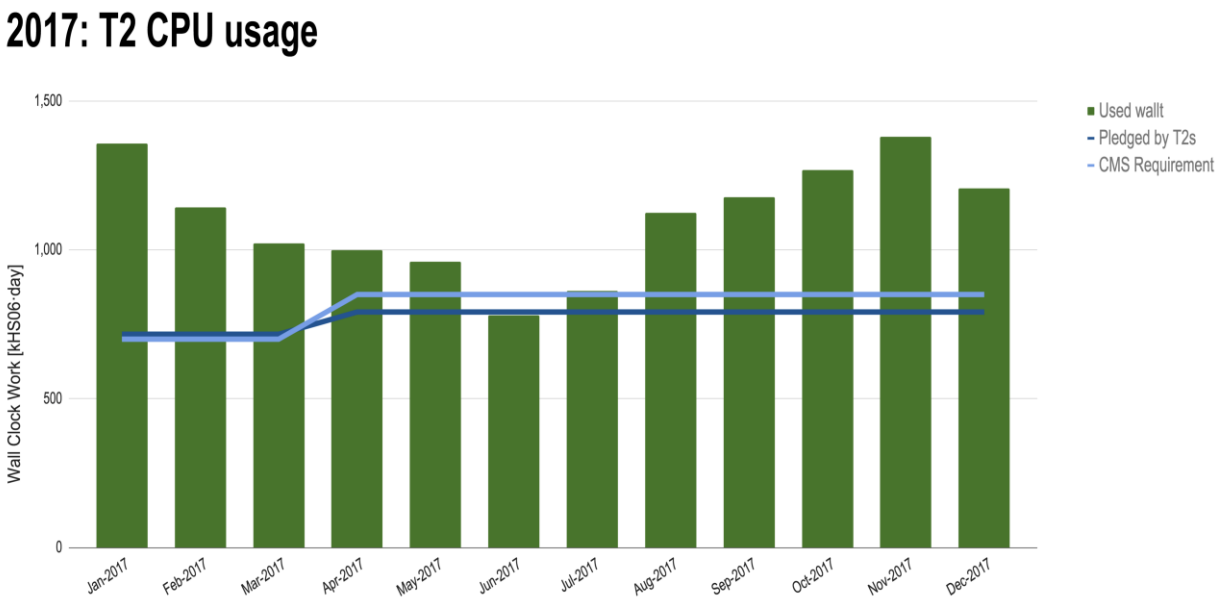

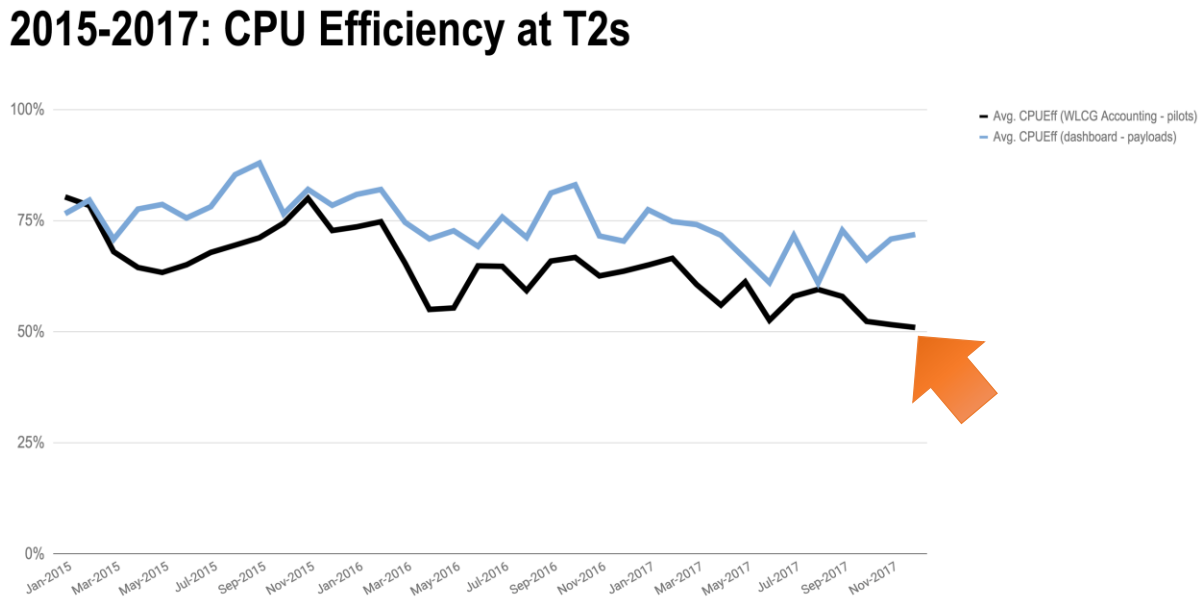

How much/many?

- Gathered monthly CPUTime and WallTime HS06·days for all of the CMS Tier-2 sites in the EGI accounting portal for 2015, 2016 and 2017 and did some trivial checks:
  - Is monthly CPUeff = 100% or > 100%? → 4.2% of months affected – 6/71 sites
    - CPUTime>0, WallTime>0 & CPUTime = Walltime?
    - CPUTime>0, WallTime>0 & CPUTime > Walltime?
  - Evaluation of CPUeff average in 2017 + stddev – detection of months with low eff. (away >3*stddev low)
    - 3 big sites detected: Pisa, Nebraska, Purdue
  - CPUTime=0, WallTime>0? → 0.95% of months affected – 12/71 sites
  - CPUTime>0, WallTime=0? → No occurrences

  - gSheet: https://docs.google.com/spreadsheets/d/1Cu9kWt22LEgFgEXSpGF05AEaDNpJG6AUXoFLwGRAmCQ/edit?usp=sharing

Examples

This view shows the accounting data in WLCG Tier2 Site BelGrid-UCL. The metric shown is CPU Efficiency, grouped by VO and Month, a custom selection of VOs are shown. No local jobs are shown

Resource Centre BelGrid-UCL — CPU Efficiency (%) by VO and Month (Custom VOs)

VO	Jan 2017	Feb 2017	Mar 2017	Apr 2017	May 2017	Jun 2017	Jul 2017	Aug 2017	Sep 2017	Oct 2017	Nov 2017	Dec 2017	Jan 2018	Total
cms	85.33%	78.22%	82.06%	92.76%	100%	100%	100%	100%	100%	100%	100%	100%	100%	98.08%
Total	85.33%	78.22%	82.06%	92.76%	100%	100%	100%	100%	100%	100%	100%	100%	100%	98.08%

1 - 1 of 1 results ‹ 1 › Number of rows per page 30

Download JSON Data / Download CSV Data

The information in the previous table is also shown in the following graph.

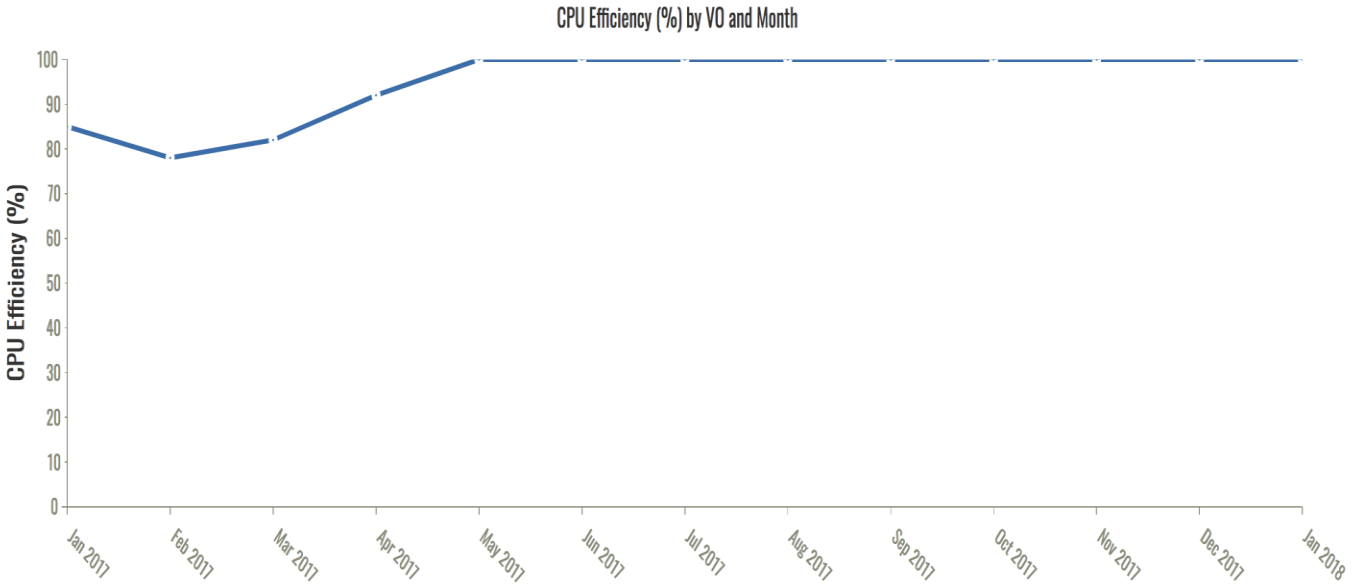

Examples

This view shows the accounting data in WLCG Tier2 Site CSCS-LCG2. The metric shown is CPU Efficiency, grouped by VO and Month, a custom selection of VOs are shown. No local jobs are shown

Resource Centre CSCS-LCG2 — CPU Efficiency (%) by VO and Month (Custom VOs)

VO	Jan 2017	Feb 2017	Mar 2017	Apr 2017	May 2017	Jun 2017	Jul 2017	Aug 2017	Sep 2017	Oct 2017	Nov 2017	Dec 2017	Jan 2018	Total
cms	98.76%	100.71%	100.35%	100.27%	99.9%	94.65%	99.39%	91.23%	93.48%	82.52%	89.28%	95.65%	95.07%	95.63%
Total	98.76%	100.71%	100.35%	100.27%	99.9%	94.65%	99.39%	91.23%	93.48%	82.52%	89.28%	95.65%	95.07%	95.63%

1 - 1 of 1 results

‹ 1 › Number of rows per page 30

Download JSON Data / Download CSV Data

The information in the previous table is also shown in the following graph.

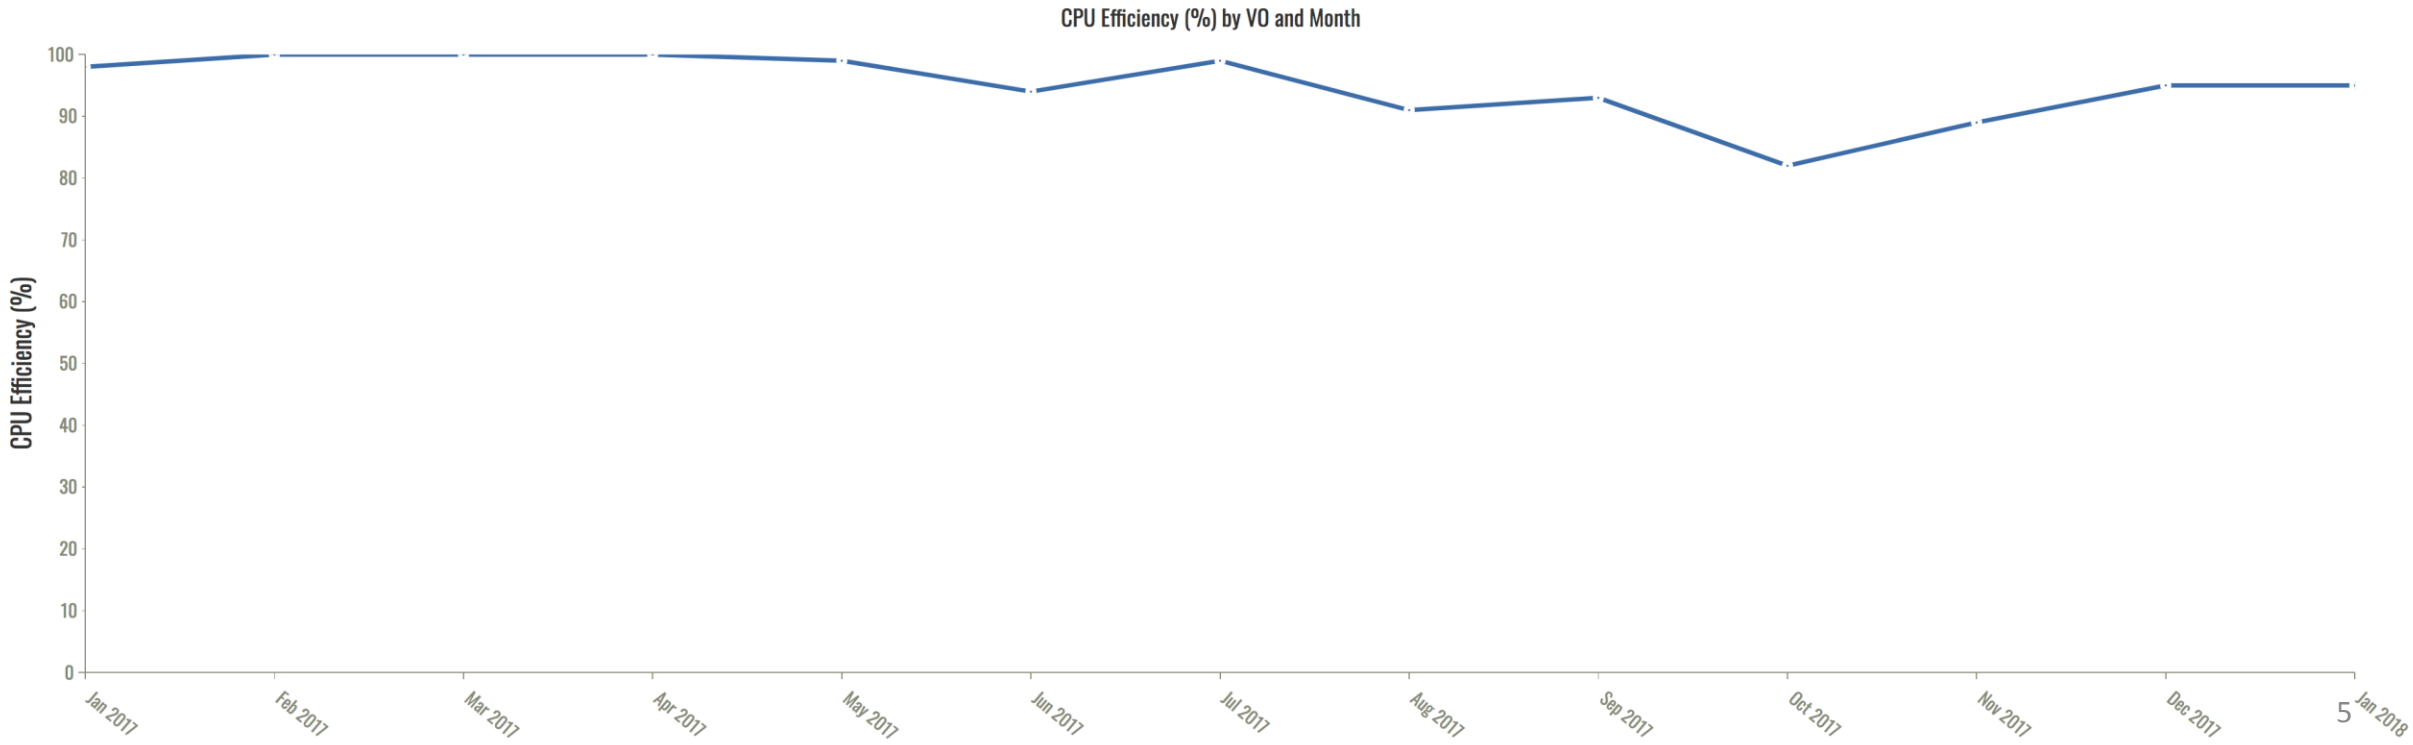

Examples

This view shows the accounting data in WLCG Tier2 Site ICM. The metric shown is CPU Efficiency, grouped by VO and Month, a custom selection of VOs are shown. No local jobs are shown

Resource Centre ICM — CPU Efficiency (%) by VO and Month (Custom VOs)

VO	Jan 2017	Feb 2017	Mar 2017	Apr 2017	May 2017	Jun 2017	Jul 2017	Aug 2017	Sep 2017	Oct 2017	Nov 2017	Dec 2017	Jan 2018	Total
cms	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	99%	99.92%
Total	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	99%	99.92%

1 - 1 of 1 results ‹ 1 › Number of rows per page 30

Download JSON Data / Download CSV Data

The information in the previous table is also shown in the following graph.

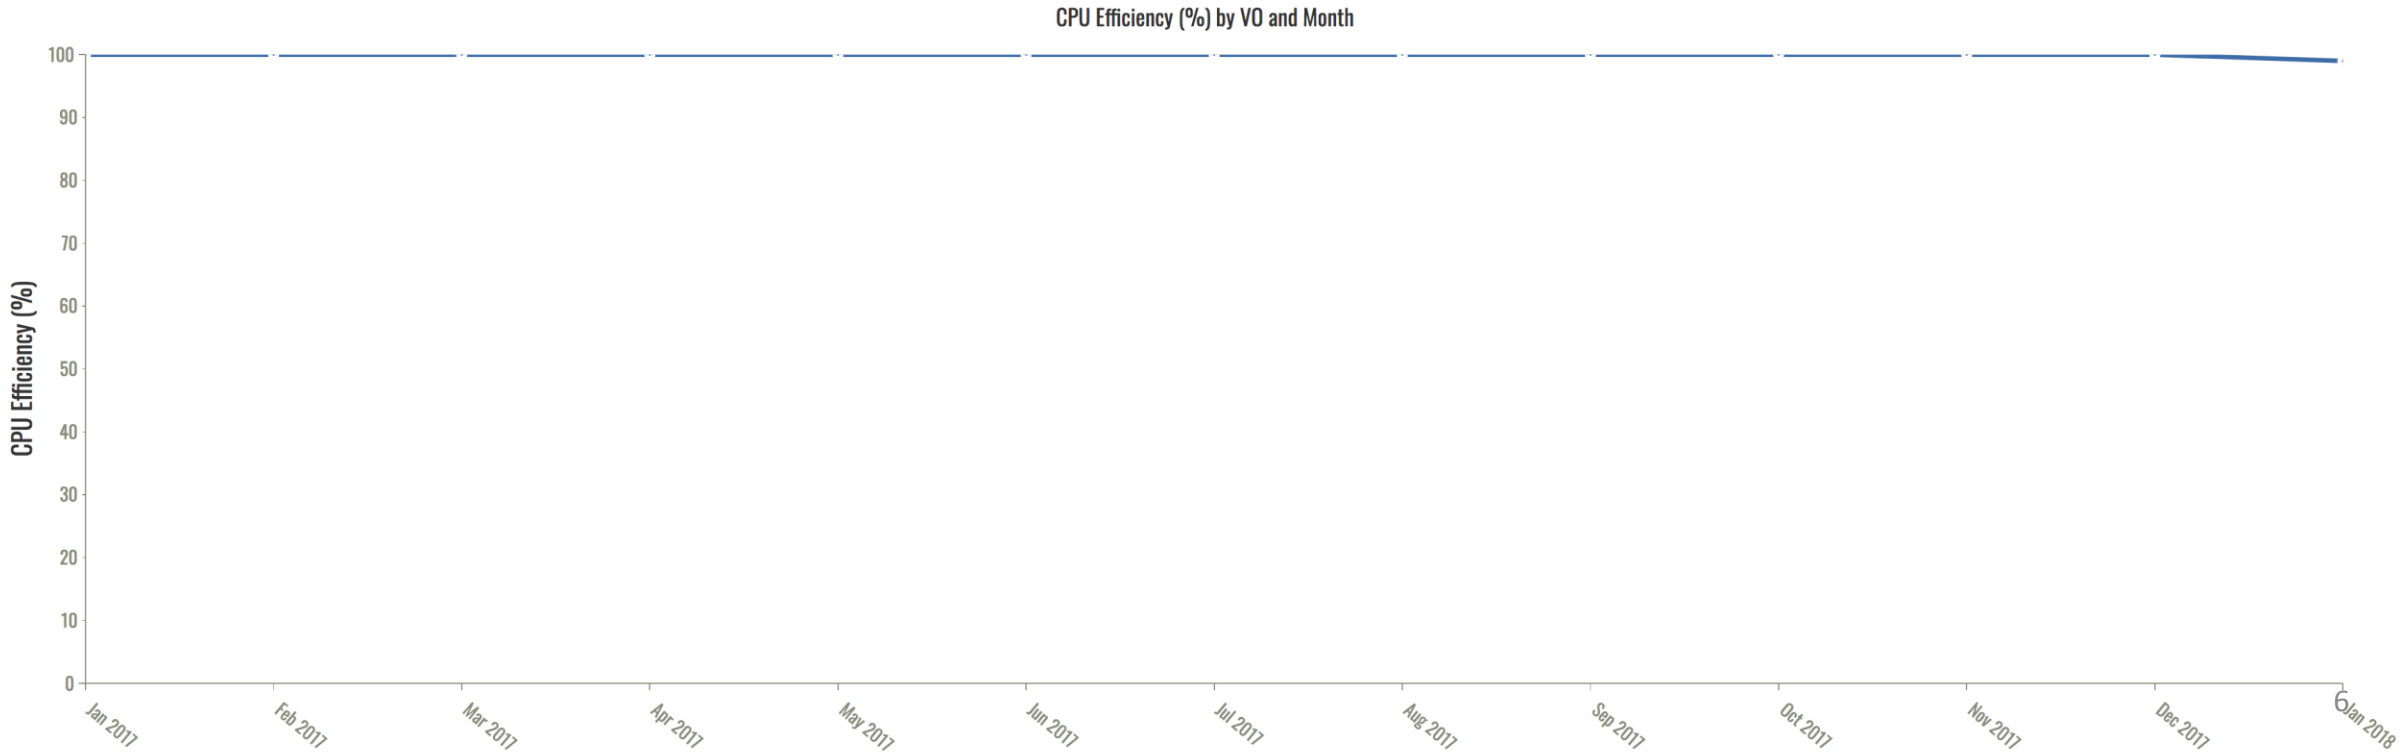

Examples

This view shows the accounting data from all Grid Sites that are classified as WLCG Tier2 in Federation T2_US_Purdue. Only production certified Sites are included. The metric shown is CPU Efficiency, grouped by Site and Month, a custom selection of VOs are shown. No local jobs are shown

Federation T2_US_Purdue — CPU Efficiency (%) by Site and Month (Custom VOs)

Site	Jan 2017	Feb 2017	Mar 2017	Apr 2017	May 2017	Jun 2017	Jul 2017	Aug 2017	Sep 2017	Oct 2017	Nov 2017	Dec 2017	Jan 2018	Total
Purdue-Carter	65.83%	68.23%	61.3%	48.08%	75.65%	0%	0%	0%	0%	0%	0%	0%	0%	61.69%
Purdue-Conte	26.31%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0.11%
Purdue-Hadoop	86.35%	80.07%	73.98%	78.47%	82.76%	76.74%	77.4%	82.37%	89.38%	84.54%	79.47%	78.33%	82.95%	81.39%
Purdue-Halstead	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Purdue-Hammer	56.55%	59.56%	54.41%	41.31%	57.67%	46.3%	50.12%	47.02%	51.54%	0%	54.68%	62.36%	0%	44.15%
Purdue-Rice	30.01%	53.59%	39.24%	33.54%	53.59%	53.9%	50.78%	56.2%	65.33%	0%	31.47%	42.26%	0%	37.58%
Total	60.15%	63.28%	56.69%	38.97%	51.45%	44.59%	47.21%	44.14%	44.38%	2.45%	43.86%	47.4%	1.3%	41.75%

1 - 6 of 6 results ‹ 1 › Number of rows per page 30

Download JSON Data / Download CSV Data

The information in the previous table is also shown in the following graph.

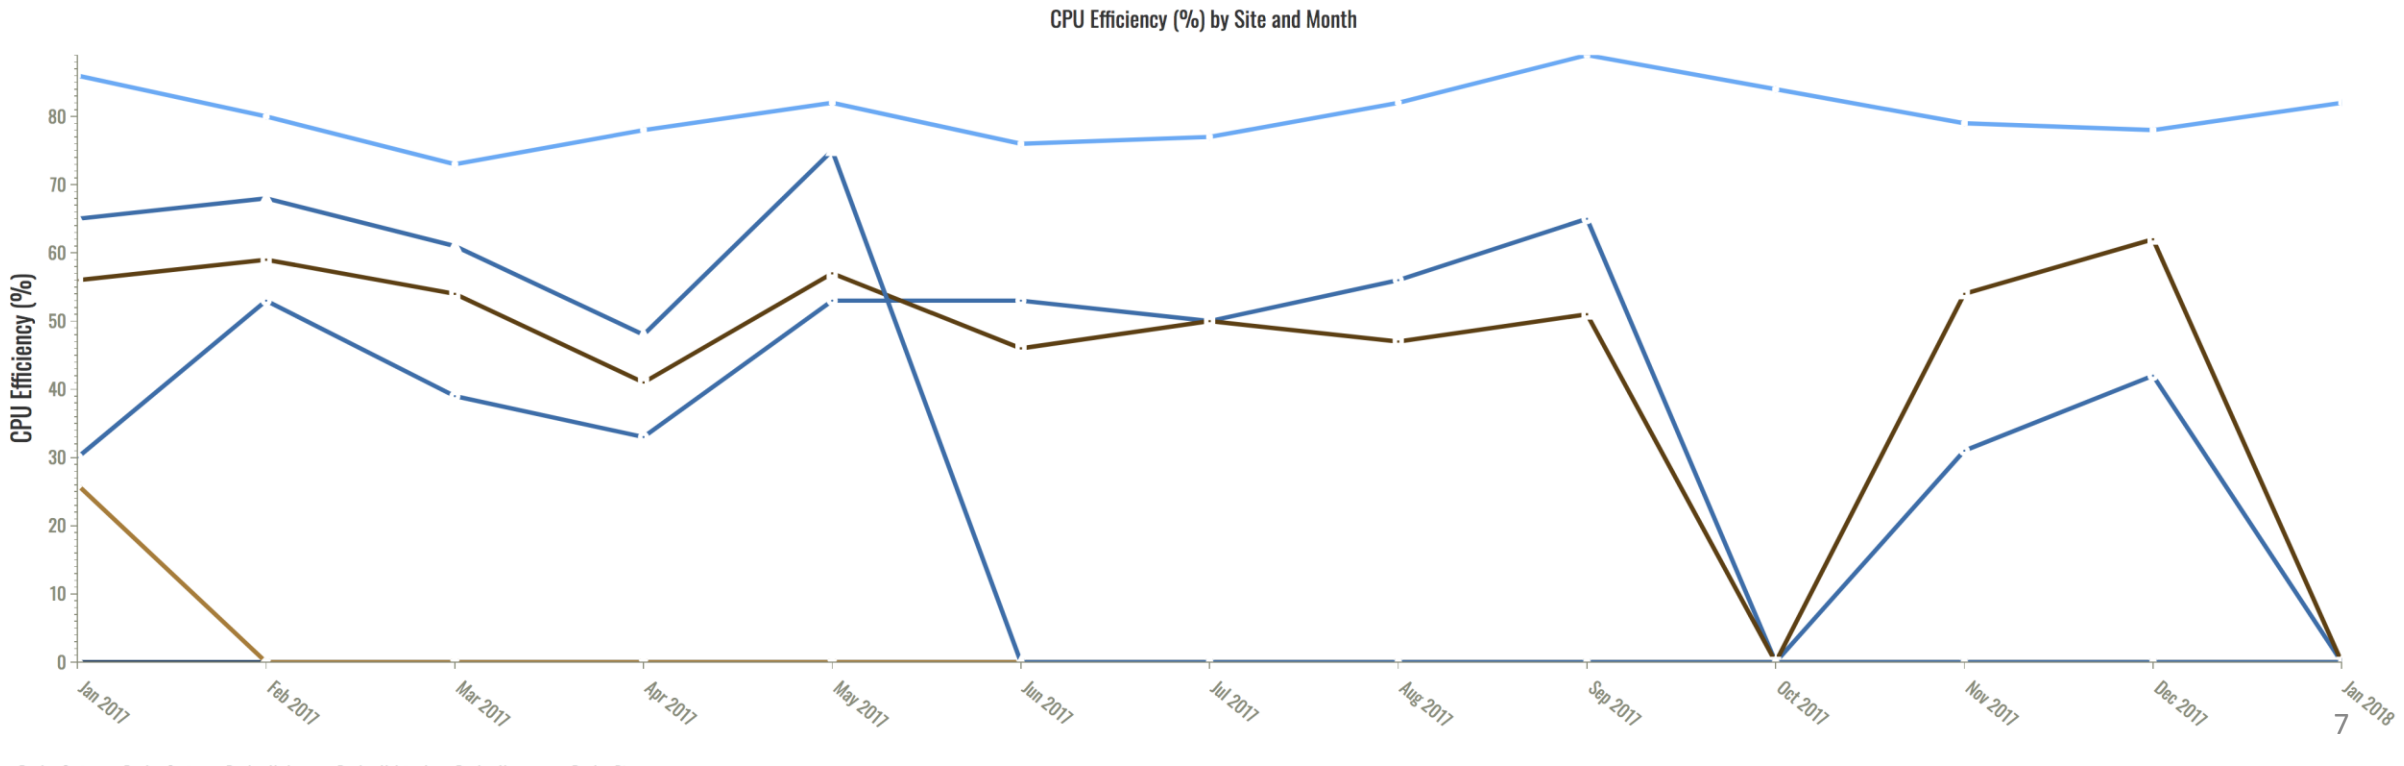

SLURM problems related

T3_UKI_LT2_QMUL:https://ggus.org/?mode=ticket_info&ticket_id=132929 **[In progress]**
<-- this site appears as a Tier-2 in the accounting.egi.eu, so another thing in the to-do list, correct the topology there!
Problem: many months in 2017 with cputime=walltime. CPU_Eff = 100%.
Response
SLURM issue
The cpu time is taken from CPUTimeRAW, and the wall time is taken from Elapsed) they are the same
wall=$(/usr/bin/sacct --allocations -P -n --format=Elapsed -j $jobid)
cpu=$(/usr/bin/sacct --allocations -P -n --format=CPUTimeRAW -j $jobid)

According to NDGF Tier-1 SLURM expert. Here are some of the interesting fields per job:
sacct -P -n --format=Elapsed,NCPUs,CPUTimeRAW,UserCPU,SystemCPU -j
And WLCG "CPU Time" is typically interpreted as UserCPU+SystemCPU.

Added J. Gordon

T2_PL_Warsaw:https://ggus.org/?mode=ticket_info&ticket_id=132922 **[In progress]**
Problem:since many months, site publishing the same time for cputime and walltime
Response
SLURM batch system
Probably incorrect sacct format was used to generate log for apelparser

T2_BR_UCL:https://ggus.org/?mode=ticket_info&ticket_id=132925 **[In progress]**
Problem: many months in 2017 with cputime=walltime. CPU_Eff = 100%.
Response
APEL parser problem for**SLURM batch system**, cputime = wall time? adding john G.

Other problems

T2_CH_CSCS:https://ggus.org/?mode=ticket_info&ticket_id=132927 **[In progress]**
Problem: many months CPU_eff>100%
Response
the submit to EGI direct records from ARC CE
Added J. Gordon
raw cpu/wall for cores=1 seems bad for at least 1 month
Asking which batch system is being used for further assistance

T2_PT_NCG_Lisbon:https://ggus.org/?mode=ticket_info&ticket_id=132928 **[In progress]**
Problem:since many months, site publishing cputime>walltime
Response
they claim ATLAS is ok. Don't know how to proceed.
Run GE batch system
Added J. Gordon

T2_GR_Ioannis:https://ggus.org/?mode=ticket_info&ticket_id=132926 **[solved]**
Problem: Feb. 2017, site was publishing cputime = 0 for all of the jobs executed in the month (!=0, walltime!=0).
Response
hacker attack in Feb. 2017. This explains cputime=0

T2_UA_KIPT:https://ggus.org/?mode=ticket_info&ticket_id=132923 **[solved]**
Problem: many months CPU_eff>100%
Response
They missed the "parallel=true" in the APEL config. Re-published.

Big problems

T2_IT_Pisa:https://ggus.org/?mode=ticket_info&ticket_id=132924 **[In progress]**
Problem:they stopped publishing values since Jul. 2017
Response
It seems that the problem start after tha installation of new server certificate.
Working on it. There is a ticket opened to APEL: https://ggus.org/?mode=ticket_info&ticket_id=132889

T2_US_Purdue:https://ggus.org/?mode=ticket_info&ticket_id=132930 **[In progress]**
Problem: since Jul. 2017, some instances of Purdue are publishing through APEL cputime=0 and walltime!=0 for the jobs run
Response
James L. following ticket.
The only CE that reported the cputime metric consistently throughout 2017 was our HTCondor based CE.
The others that stopped reporting are all PBS based.
Re-publishing

T2_US_Nebraska:https://ggus.org/?mode=ticket_info&ticket_id=132931 **[In progress]**
Problem: efficiencies observed at the EGI portal are very low for Nebraska.
Response
There is a bug in HTCondor involving the interplay of docker and condor where the CPU hours are not reported correctly:
https://htcondor-wiki.cs.wisc.edu/index.cgi/tktview?tn=6173

These numbers are lost/unrecoverable. CMS will have do help us move forward with establishing reasonable numbers.
(As of now we're not clear on whether the bug is corrected in the latest Condor, we're testing.)

CMS request to WLCG Accounting TF

- Some problems in the accounting which were not (yet) noticed by this group
- Could these simple checks be added in SSB and/or be checked from time to time?
  - We need better alarming system for these simple and trivial accounting failures
  - Other VOs for sure affected
- CMS understands that the responsibility to certify, validate and follow-up WLCG accounting issues is on WLCG (aka this group)
  - We are still working in more and better ways to catch these issues, but any help is welcome!